10702

GRANT

IN-55-CR

140850

P- 20

Wide-Bandwidth High-Resolution Search for Extraterrestrial Intelligence

Semiannual Status Report

15 Dec 1991 -- 15 June 1992

Paul Horowitz, Principal Investigator

Harvard University
Cambridge, MA
02138
June 15, 1992

Grant Number NAGW-2872

1. INTRODUCTION

This interim report summarizes research accomplished during the initial 6-month period of the grant. Toward the end of this period we enrolled a graduate student, Darren Leigh, who finished his MSEE at MIT in June, 1992. Thus during the period covered by this report the active personnel included the PI, two Harvard undergraduates (Neil Hendin and Greg Galperin, both part-time), and minor participation by Mr. Leigh, who has now commenced full-time work on the project. For the summer months we have been joined by another undergraduate, Derek Bass.

2. RESEARCH ACCOMPLISHED

2.1 Antenna Configuration

Following conversations with Prof. David Staelin (MIT) and with Dr. Michael Davis (Arecibo), and after considerable work with 5 years of data from the ongoing META search apparatus (see *Publications*, below), we have concluded that the problem of terrestrial interference is sufficiently severe that it is worthwhile including a "terrestrial" feedhorn (and spectrometer), in parallel with the East and West feedhorns originally planned. Each feed will drive an identical 80-megachannel spectrometer, based on parallel replication of the Berkeley Serendip-III architecture, as originally proposed. The terrestrial feed will probably consist of an azimuthally-symmetric low-gain (+3dBi) pattern; simultaneous detection of a signal in the terrestrial beam and one of the (high-gain) sky beams results in immediate veto.

Another suggestion from Staelin is that the two sky beams should include an area of overlap. This mandates a small phased array of horns, which we are currently designing. A promising design uses 7 horns and three low-noise preamps: The central horn is buffered and phased with each of the two buffered passively combined outer sets of three horns. The advantage of overlapping beams is the observable "handoff" of a continuing signal; we expect to make the decision as to whether this feature is worth the extra effort of a phased array during the current 6-month period of the grant, and move forward with construction on the method of choice.

2.2 Downconverter

Considerable progress took place in this area. The channelizing downconverter requires an array of 20 local oscillators (LO's), equally spaced in frequency across the IF bandwidth, and preferably phase-locked to the master station clock. We looked at several technologies for this synthesis, namely *i)* direct digital synthesis (DDS), *ii)* an array of fixed crystal oscillators, *iii)* a comb generator, driven at the channel spacing frequency, driving a set of narrow filters and buffers, and *iv)* an array of phase-locked-loops (PLL's). We concluded that DDS is too expensive, and prone to spectral impurities (spurs and harmonics); crystal oscillators, though delivering quite good signal purity, are difficult to obtain in small quantities at specific frequencies (in response to our specifications, most vendors bid "no bid"!); the comb generator approach requires an elaborate set of high-Q (and potentially drift-prone) filters at VHF; and thus the PLL approach seemed best.

PRECEDING PAGE BLANK NOT FILMED

A good PLL requires a good open-loop oscillator, so we began by exploring various designs of varactor-tuned VHF oscillators. In particular, we look at a Meissner circuit (transformer coupled, JFET), an IC circuit (Motorola MC1648), and a JFET Clapp oscillator. The latter was the best, delivering both open-loop and locked spectral purity better than our Hewlett Packard 3325A HF synthesizer. We measured both wideband (to 100 MHz) and close-in (down to 1.5 Hz full-span) spectra, demonstrating better-than-required spectral purity (see attached spectra, Figures 1-4). To measure close-in noise, we mixed the VHF LO with the synthesizer, both referenced to the same master crystal, then lowpass filtered the beat note (7-pole Chebyshev, $f_o=1$ MHz) and analyzed it with a low-frequency (100 kHz) FFT analyzer. The resulting spectra are, of course, the convolution of the separate spectra, and set a worst-case limit of spectral purity that is far better than we need. We used a Schottky-diode amplitude feedback arrangement, to prevent JFET clipping or gate conduction, a common condition in casually designed JFET oscillators; the result is an extraordinarily pure sinusoidal oscillation, compared with the usual Q-robbing clamped waveform (see Figure 5).

It was interesting to note that the VHF oscillator IC (MC1648), which claimed "high spectral purity," was by far the worst of the three oscillators, being some 10-30dB poorer than the others (Figure 6; note carrier level is reduced 15dB compared with Figures 1-3).

We next closed the loop, using initially the 74HC4046 phase detector; it performed poorly, displaying serious phase jitter. We traced the problem to the zero-phase "dead band," a frequent problem in frequency synthesis, for which there are improved phase detectors. We tried the MC145151, a single-modulus PLL synthesis chip with parallel-input modulus selection; it worked well, but was too inflexible in choice of modulus, and also occupied a large package. Our final choice was the MC145170, a superb PLL chip, in a small 16-pin DIP, with no deadbanding problem, and a good price ($6.90). It requires serial downloading, via an interesting 3-pin serial port; initially we used a 22V10 state-machine PAL to accomplish the download, but finally settled on an 87C751 microcontroller to do the job, since it can download a large array of PLL's as easily as it can do a single PLL. The schematic (Figure 7) shows the final circuit configuration of both JFET oscillator, PLL, and microcontroller downloading arrangement.

After PC breadboarding and testing, we layed out a PC board with 10 VHF LO's; two such boards are needed for the full BETA system. The board is two-sided, with a ground plane to ensure clean signals and good isolation between oscillators (see Figure 8). The boards have been manufactured, and are awaiting a final decision on channel spacing before being fully stuffed. We made some extra boards, at the request of the Berkeley group, which may use them in its channelizing filter bank.

The LO array drives the mixer/filter/ADC subsystem, on which we made good progress also. We began by designing our own anti-alias lowpass filters, because commercial filters were quoted at $50 per filter (and 40 are needed per feedhorn!). We built 5-pole Butterworth and 7-pole Chebyshev versions (the latter in both 0.01dB and 0.1dB passband ripple configurations), all with a 1.0 MHz cutoff (anticipating 2 MHz of real-time bandwidth per spectrometer path). Figure 9 shows a linear/linear sweep, dc to 2 MHz; the 7-pole, 0.1dB filter seems optimal, and

costs $8 per copy. It uses slug-tuned inductors and pairs of mica capacitors, in an easily manufactured arrangement.

The next step is the ADC. We chose the TRW TMC1175 half-flash 8-bit ADC, which costs less than $10 and converts up to 20 MSPS (our application requires only 2 MSPS). From our experience with META (and after comparing notes with Berkeley), we decided that open-loop trimming of the dc offset is marginal in systems with large FFTs, because even a fractional LSB dc offset puts a "coherent" signal into the dc channel that can produce numeric overflow in a system scaled for noiselike signals. For example, 1 LSB of dc offset produces 2^{22} unscaled counts in the dc channel of a 4 megachannel FFT, which might have something like 14 bit scales to prevent (incoherent) word growth with the expected noiselike signals. That amounts to 8 bits of coherent signal after scaling. In the ongoing META search, we carefully trimmed the dc offset, while watching the FFT's dc channel; it was a critical adjustment, with potential for catastrophic drift.

Our solution here is a unique "auto-zero" circuit, which to our knowledge has not been used elsewhere. The idea is to accumulate a running sum of the (signed) digitized outputs of the ADC, then use the output (a digitally integrated error signal) to drive a DAC that trims the analog input offset. Of course, this has to be done for both the I and Q mixer channels. The schematic diagram in Figure 10 shows our circuit solution. The accumulation is performed in a 44-pin "super-PAL" device from AMD (their Mach 110, with 32 macrocells), suitably programmed; we managed to squeeze a 16-bit adder/accumulator into a single part, with a remarkable 96% product-term utilization. The DAC08 is a very fast (85ns) 8-bit current output DAC that costs less than $2, and is configured to provide about 4 LSB's of trim over its full range. In this circuit the DAC is tied to the top 8 bits of the accumulator, which carries out of the bottom 8 bits, where the input quantities are being added.

The autozero successfully keeps the dc signal near zero, in the process introducing a low-frequency rolloff in the converter response. That occurs because the input signal is ac coupled, with the autozero trimming the dc input to the ADC just downstream of the blocking capacitor. Our circuit puts the rolloff at about 600 Hz, which is desirable in any case because of 1/f type noise in the vicinity of the dc channel. In fact, in our current system we let the FFT cope with large amounts of dc and 1/f spectral components, then postprocess the spectrum by setting those components to zero! The autozero accomplishes the same thing, without straining the FFT numeric headroom.

There's an interesting serendipity in this circuit, in that the high frequency cutoff of the autozero feedback signal is precisely the same as the low frequency cutoff of the (blocked) input signal; that is because both are determined by the blocking capacitor in combination with the dc resistance seen at the ADC input. We demonstrated this by injecting a logarithmic frequency sweep (dc to 10 kHz), and plotting (top to bottom traces, in the top oscilloscope photo of Figure 11) the input signal, the signal reaching the input to the ADC, and the autozero feedback signal.

The autozero circuit performs very well. The remaining oscilloscope photos in Figure 11 show the dynamic performance, with a half-scale wideband noise input (top trace, 1V/div). The middle photo shows the DAC output, with the loop open; the sawtooth is the result of a few

LSB's of average dc offset, accumulated to a digital ramp. The lower photo shows the result when the loop is closed; now the autozero holds the input trimmed, with only a small portion of the DAC range exercised, even with a large input noise signal, of the sort routinely expected.

In our prototype, the autozero is easy to use: Just trim the manual offset adjustment to bring the DAC output to the middle of its range. Small changes in trim simply cause the DAC output to shift up and down, when the autozero loop is closed, whereas the trim causes changes in *sawtooth wave frequency* when the loop is open. In the latter case it is nearly impossible to make the DAC output stationary; furthermore, even if accurate open-loop trim can be achieved, it is lost when the input *amplitude* changes, because the overall dc trim depends on the symmetry of ADC decision levels, as averaged by the input signal. I.e., the ADC dc trim is a dynamic quantity, hence untrimmable *in principle* without an autozero!

We are completing final testing of the autozero circuit, and checking out an alternative ADC part (the TI TLC5502-5), which is a comparably-priced *full-flash* 8-bit converter, compared with the half-flash TRW part. When tests are complete, we will reduce the circuit to PCB pattern (this time using CAD/CAM), and have boards manufactured. Together with the now-completed LO array, this will complete the downconverter task.

2.3 FFT Array

We began with a survey of five commercially available FFT ASIC's: The Array Microsystems 66111, Austek A41102, LSI Logic L64280, Plessey 16510, and TRW TMC2310-1. In each case we studied data sheets and made preliminary comparisons, then proceeded to block out implementations. Some of these chips include automatic block floating point scaling, a good feature for individual FFT's, but an awkward arrangement when a long 1-dimensional FFT is performed as a succession of shorter row and column FFT's, as we plan to do. That is because, in a transform of N*M points, one must interpose a complex multiply (known as a "twiddle factor") between the N row and M column FFT's during the "corner turn"; thus, if the FFT chip has performed a block floating point scaling, each row, say, will in general have a different exponent, as the chip scales each short transform. Now the column transform becomes extremely awkward, requiring one to find the largest row exponent, then barrel-shift all other rows to the same exponent.

Taking account of this sort of problem, and considering overall bandwidth issues, we eliminated three of the FFT chip sets, leaving the Austek and LSI Logic as finalists. The LSI excels in precision, but is expensive and requires special 44-bit wide video shift registers for intermediate storage; it also mandates an awkward recirculating configuration for efficient low-bandwidth applications, having been designed for wideband pipelined use. Among other things, this requires that the clock rate be lowered, and that input and output data be handled in rapid bursts of low duty cycle, with consequent poor use of memory resources. By contrast, the Austek chip is ideal for continuous-flow real-time data streams, requires no special memory, and has great flexibility in terms of bit scaling, word width, transform length, normal vs bit reversed sequence, and use of internal multiplier and data switches. In addition, its documentation is unparalleled. On the down side, it (like the others) is a sole-sourced product, with future supply not assured.

Based on the above considerations, and given Berkeley's choice of the Austek chip and negotiation of an excellent price, we placed an order for 300 pieces, enough to build two complete BETA systems. In part this was a conservative response to Austek's decision to discontinue this excellent part, allowing us one "lifetime buy."

The order was placed at the end of December, 1991. While waiting for delivery we began simulations, to determine the needed word precision (16 vs 20 vs 24 bits, fixed point), and also to determine optimum twiddle factor ROM organization. We began by favoring 20 bit arithmetic, and a 16-bit by 256K ROM quarter-sine lookup table (which provides 17-bit precision after appending the sign). We also devised an efficient 3-chip 4-megachannel FFT configuration, taking advantage of the Austek chip's ability to perform normal or bit-reversed transforms (see Figure 12).

It now appears that 16-bit arithmetic is adequate, though not extravagant, for a 4M FFT with noiselike input signal. Preliminary data from Serendip-III supports this conclusion. We are completing simulations, and will make the final decision shortly.

On the chip supply issue: Austek has indeed discontinued the part, with the silicon for our order now fabricated but not packaged. During this period the manufacturing rights to the chip, granted to Austek by the Australian government's CSIRO, expired, and were sold to a reputable Australian silicon foundry known as AWA. The latter will be completing packaging and delivery of the ordered chips, probably in October, 1992, at the negotiated price. Although the delay is not good news, the continuity of future supply is very good news. We will keep busy in the meantime with further simulations, PC board layout, and work on other aspects of the system.

2.4 DSP Backend Array

The FFT array is followed by an array of dedicated digital signal processing (DSP) chips, which sift through the spectra looking for i) new peaks that exceed background by a significant amount, and ii) the progress of previously flagged frequencies of interest; see the original proposal for a detailed description.

We have been looking at several choices for the DSP array, including the Motorola 56002, 96002, Intel 80960, and Star chip. The fixed-point 56002 has now been eliminated, because its address space is inadequate for the 1MB of memory it shares with the following "LAN controller" array. Likewise, the Star chip, with its multiple processors, is too complex. The Intel 80960 has recently become affordable (less than $200), and may be the best choice; however, the 96002 has the advantage of Motorola's pleasant assembly language and traditionally clean architecture, as well as having a close relationship to the 68030 processor that we favor for the LAN controller array. On the down side, it includes unneeded floating point capability.

Whatever the choice of DSP, we have decided to perform the *baseline* calculation in a dedicated hardware accumulator, using a pair of memories of length L/2 and an adder/subtractor, so that we calculate the average baseline centered on a sliding window of length L centered on the current data point. There are various choices of implementation, including FIFO or SRAM for the

buffers, and FRGAs, PALs, or DSPs for the accumulator. In fact, the comparison of current data point with baseline multiple can also be done in hardware (rather than DSP).

An interesting issue arises on the baseline computation. To do accurate signal statistics, the baseline must be calculated from *power*, computed from the I and Q amplitudes that come from the FFT. But, since the final FFT output is in bit-reversed order, the results must be stored in a "corner-turn" memory, of size 4M x wordlength. Thus, if we use a multiplier and adder to compute power, the result requires 32 bits (from 16-bit amplitudes) for corner-turn storage, thus using 16MB of RAM (or 40% of the total RAM required). A smarter approach is to store a *representation* of the power (e.g., mu-law, or modulus) in the corner-turn memory, thus halving the memory size (16 bits of compressed power, say, derived from a complex pair of 16-bit amplitudes). A good way to compute the power representation is to use a 2-step ROM approach, with I and Q indexing into a table of compressed power representation; the baseline averaging circuit, reading out of the corner turn memory, then inverts the compression to recover true power, before accumulating the average baseline. We like this method, and are working on simulations, prior to reducing it to practice.

2.5 Backend & Workstation

We have just taken delivery on our Unix workstation (SPARCstation 2), essential for serious simulations as well as for actual experiment control. It is in use around the clock! For the crucial task of backend control of the experiment, we have decided to use thin-wire Ethernet between the workstation and the LAN controller array. We are now actively building the backend software and user interface, first by building a "data-maker" that simulates the LAN controller array, complete with signals and noise. Then we will write the backend user interface and other code, allowing it to interact with this simulator on the Ethernet, exactly as in real operation. This has the advantage of letting us contend with real network issues (communication protocols, latency, congestion, etc.), as well as writing code that will work when the real DSP and LAN-controller array is substituted for the simulator.

3. NEXT STEPS

As described above, we have expanded the scope of work somewhat, by including a third spectrometer to service a terrestrial feed, and by exploring the use of phasing to produce overlapped sky beam lobes. The channel count and instantaneous bandwidth has also been increased, to 240 megachannels and 120 MHz (from the proposed 100 megachannels and 60 MHz, respectively). Though bold, we believe these changes to be important, and necessary.

The delay in FFT chip delivery is a disappointment, but probably will not seriously affect startup of the system, given other tasks awaiting completion. Overall, we are satisfied with progress to date.

4. PUBLICATIONS

During the period of this report we have submitted two publications:

Horowitz, P. *Five Years of META: Results of The Planetary Society's Search at Harvard.* Presented at the USA-USSR Joint Conference on the Search for ExtraTerrestrial Intelligent Life, August 5-9, 1991, Santa Cruz, CA (Sponsored by the US National Academy of Sciences, Academy of Sciences of the USSR, and the Life Sciences Division, NASA Headquarters, with support from the SETI Institute and the University of California, Santa Cruz). Manuscript submitted for publication in the Conference Proceedings.

Horowitz, P., and Sagan, C. *Five Years of Project META: An All-Sky Narrowband Radio Search for Extraterrestrial Intelligence.* Manuscript submitted to *Nature.*

Both describe work completed before the current grant became effective, but actual manuscript preparation took place during the period of the grant, and describe plans to be carried out during the grant period.
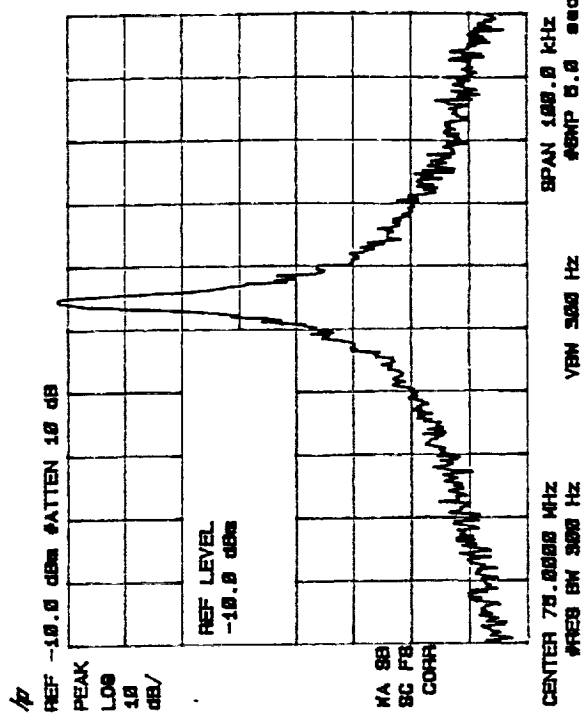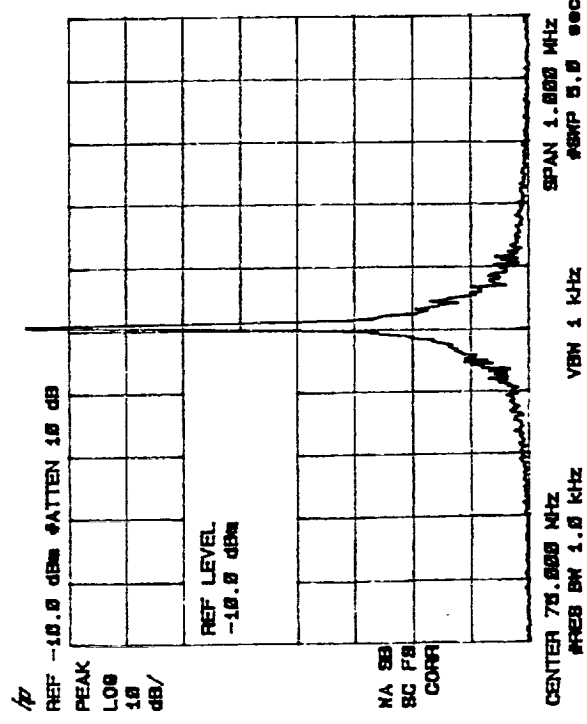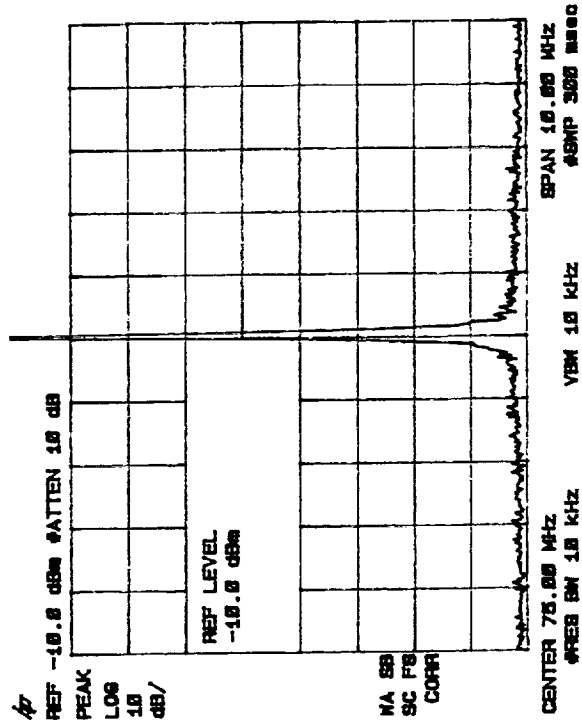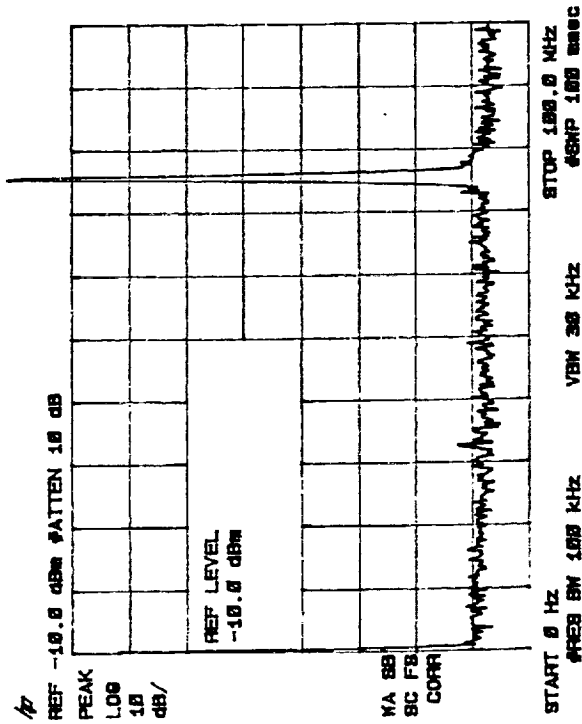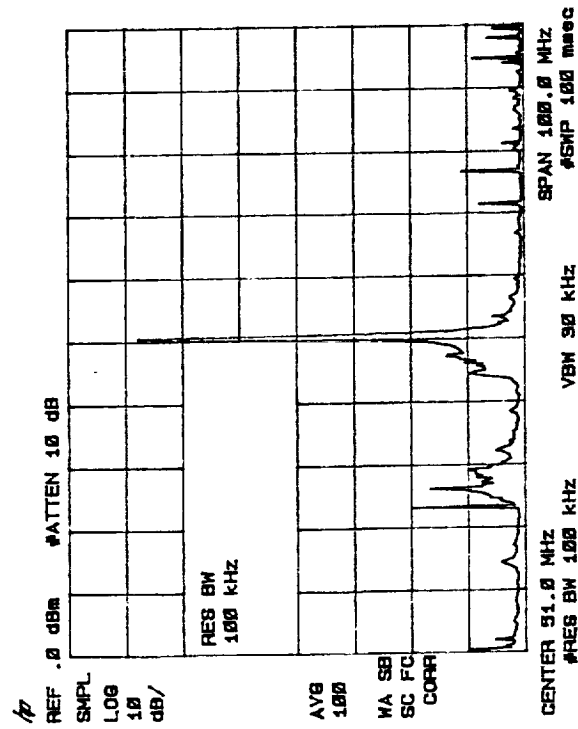
Figure 1. Broadband noise performance of LC Clapp oscillator, operating open-loop at 75 MHz. Spans are 100 MHz, 10 MHz, 1 MHz, and 100 kHz; 10dB per vertical division. Carrier level is +5dBm (15dB over top division on screen).
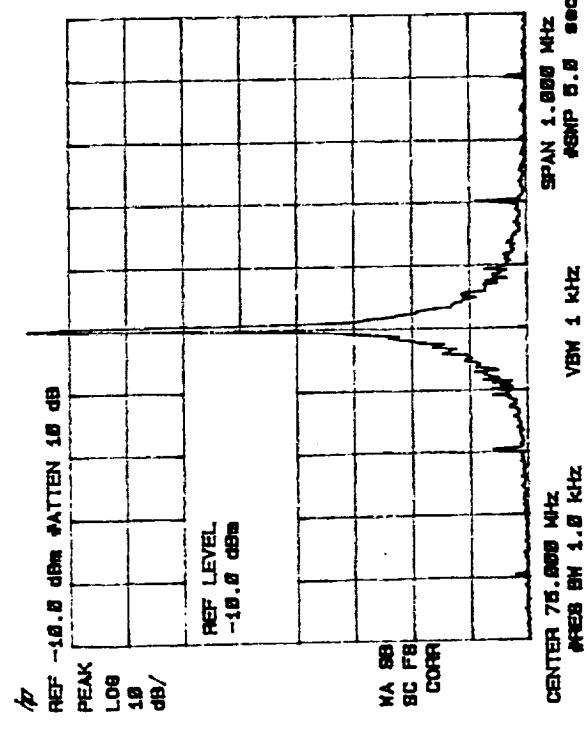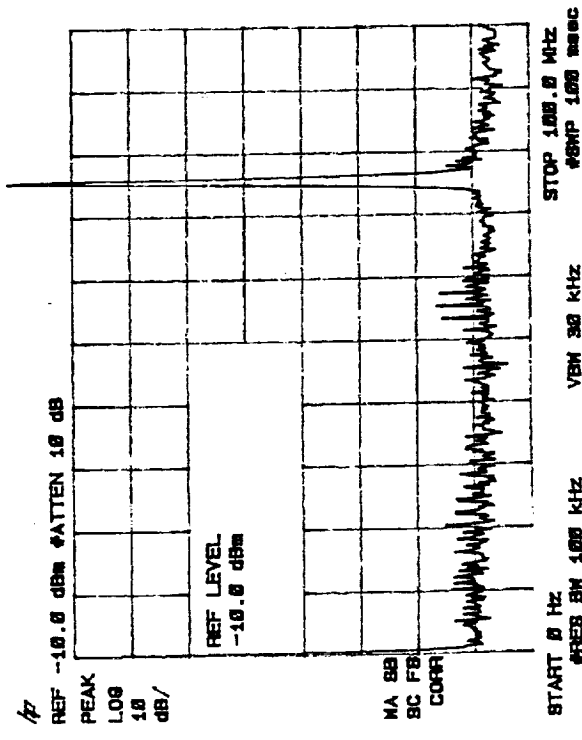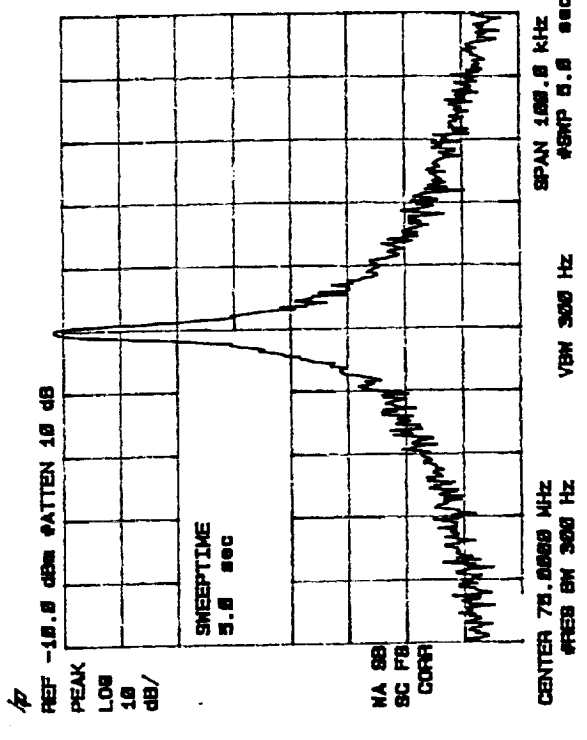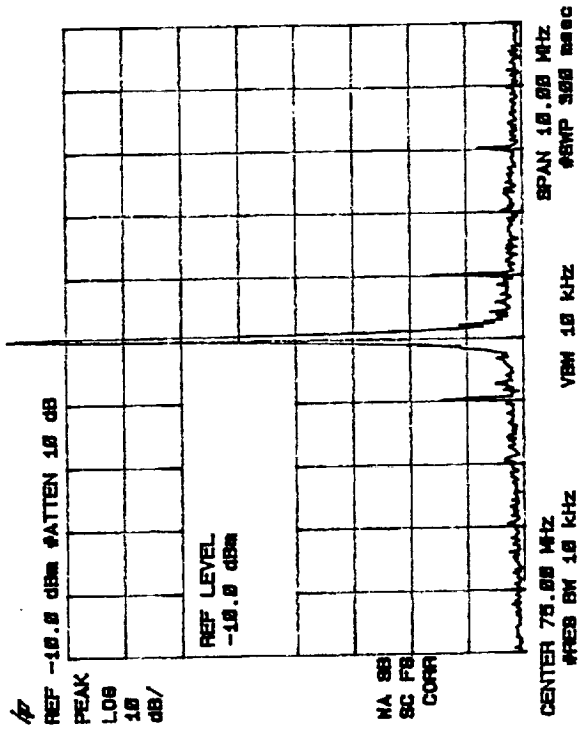
Figure 2. Same as Figure 1, except closed-loop. Note low-level spurs at 1 MHz offset (input reference frequency) and at 200 kHz offset (loop reference frequency).
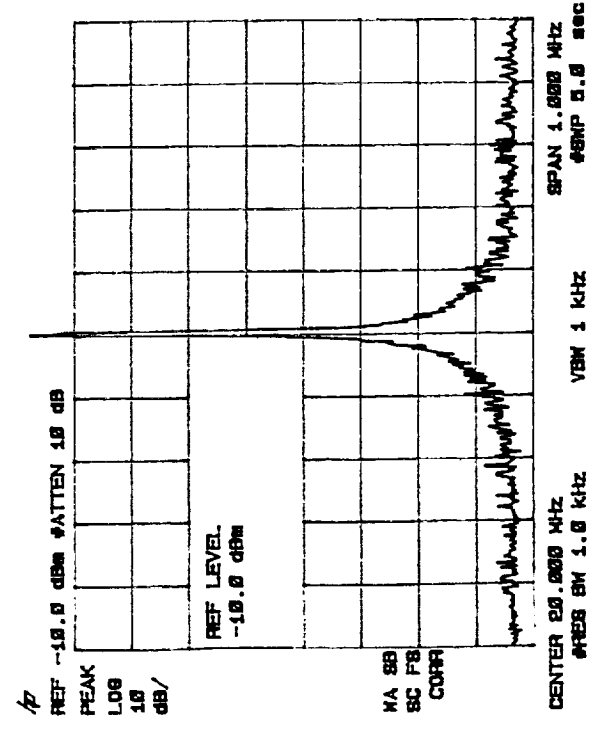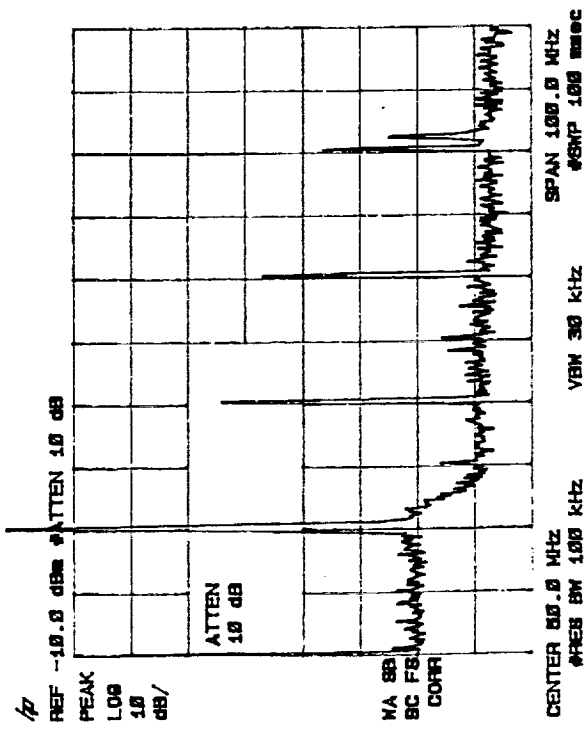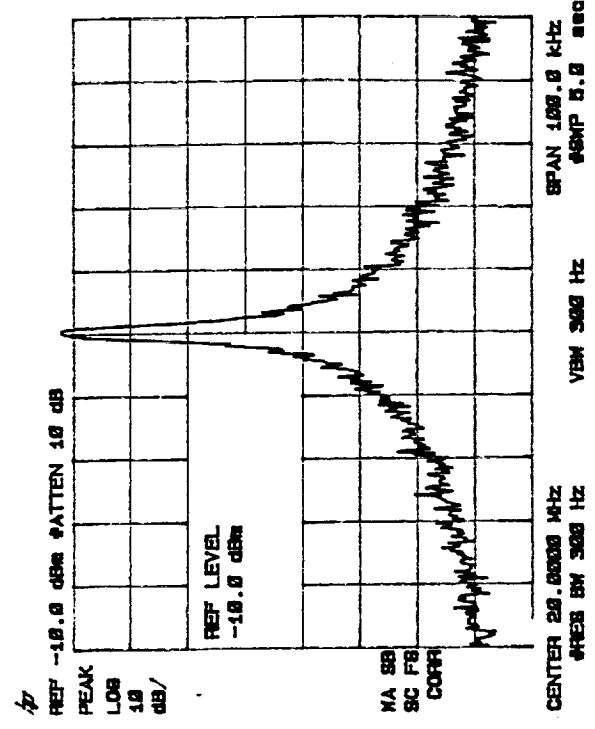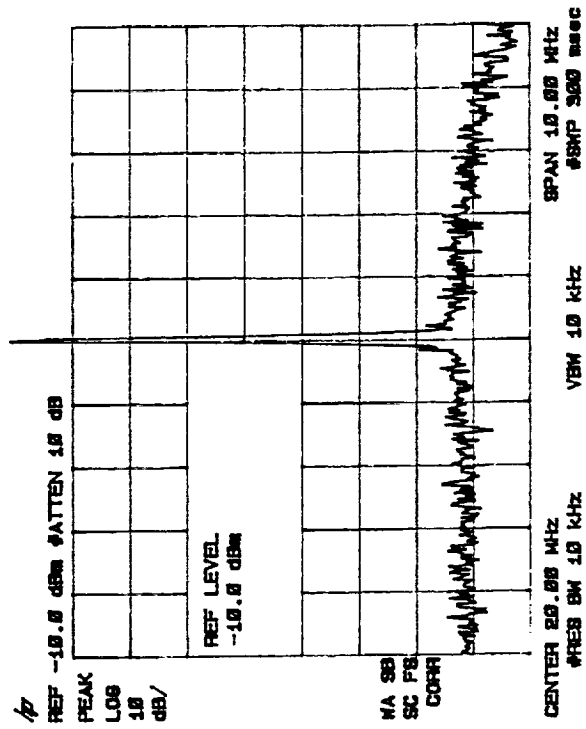
Figure 3. Broadband noise performance of HP 3325A synthesizer at 20 MHz, for comparison with LC Clapp. Same parameters as Figures 1 and 2.
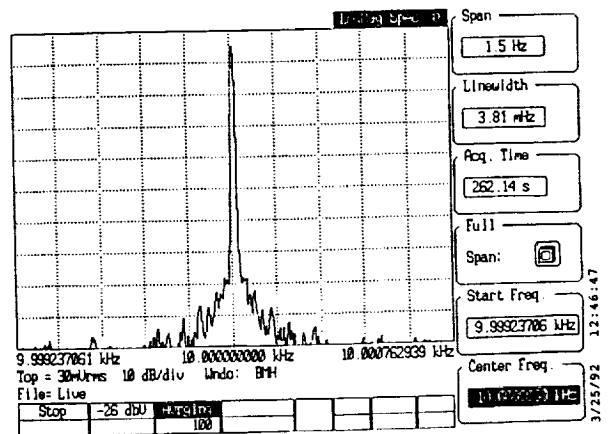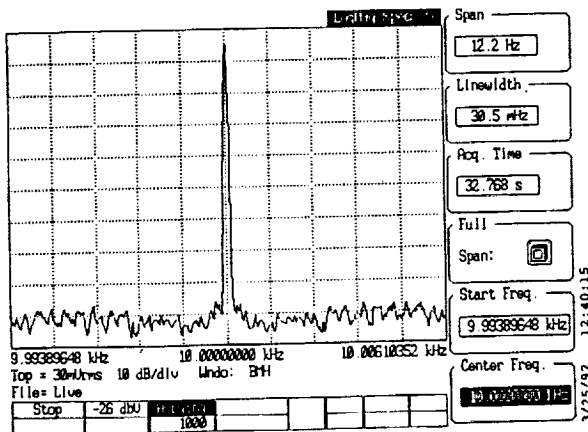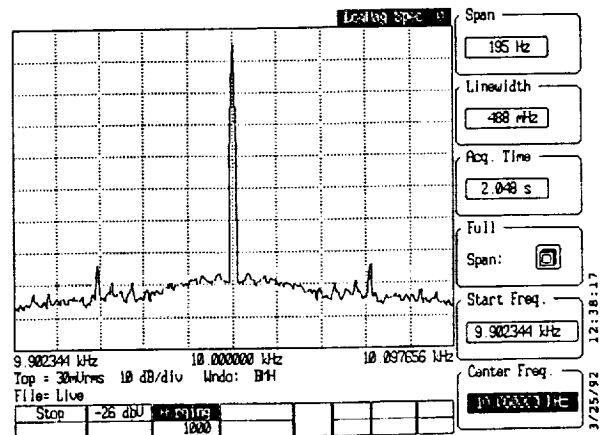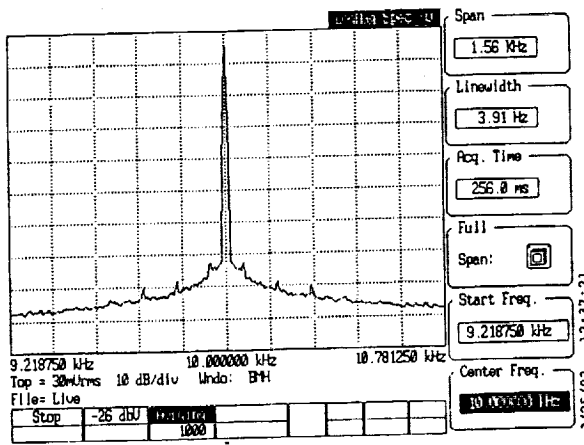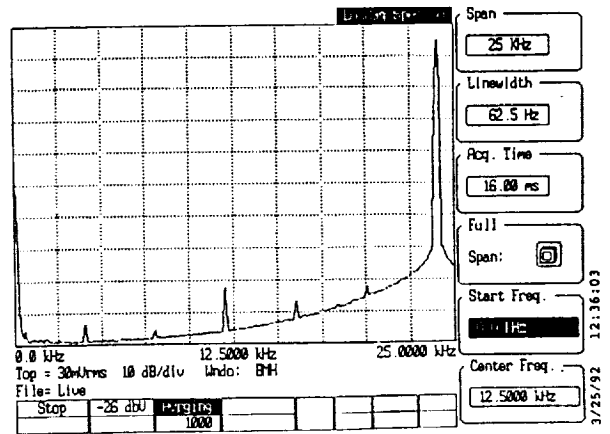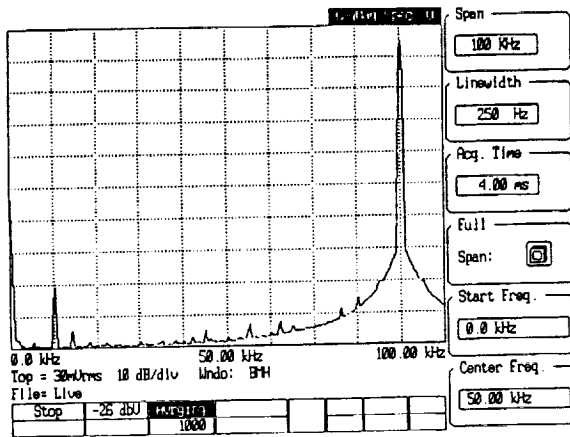
Figure 4. Close-in noise performance of closed-loop LC Clapp oscillator, programmed to 75 MHz. Spans are 100 kHz, 25 kHz, 1.5 kHz, 200 Hz, 12 Hz, and 1.5 Hz. Spectra taken with FFT analyzer, from lowpass filtered output of mixer driven by HP 3325A synthesizer. LO offset by 10 kHz, except for widest spans, where offset is 90 kHz and 24 kHz.
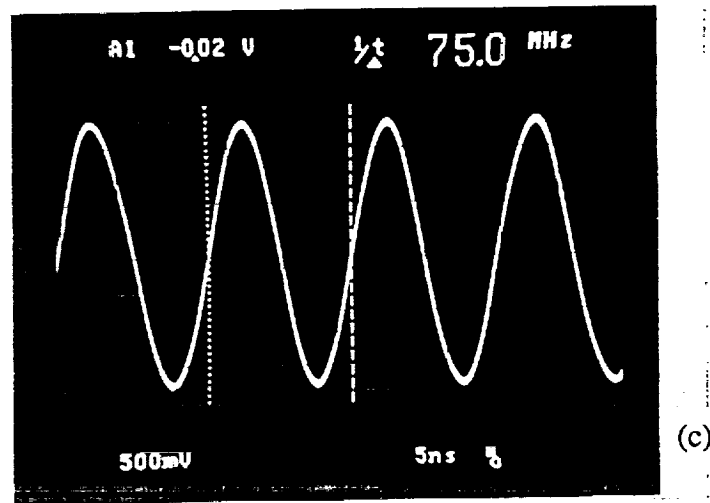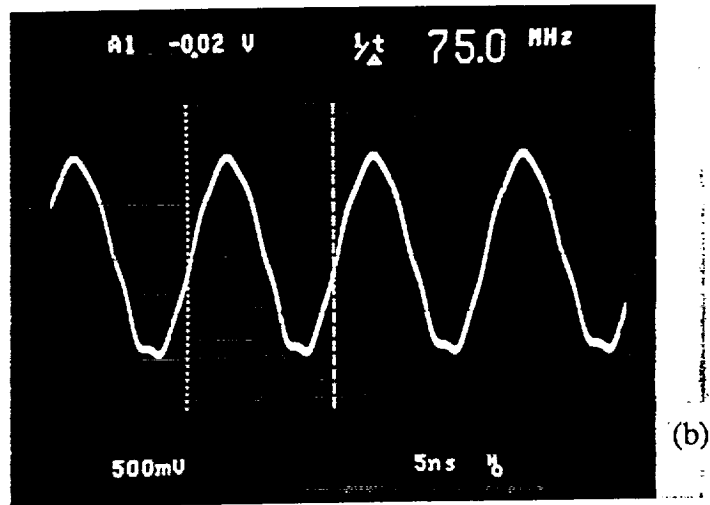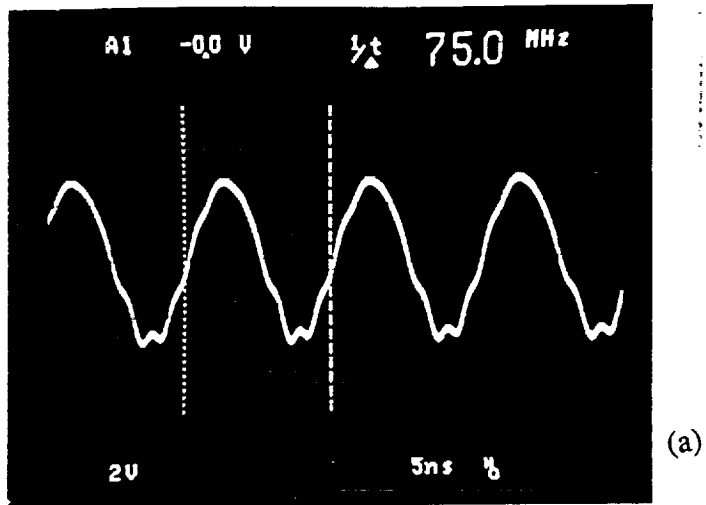
Figure 5. LC Clapp oscillator output waveforms, on Tektronix 2465B analog oscilloscope. (a) without amplitude control; (b) with back-to-back Schottky clamp diodes across output; (c) with slow feedback to gate (final circuit).
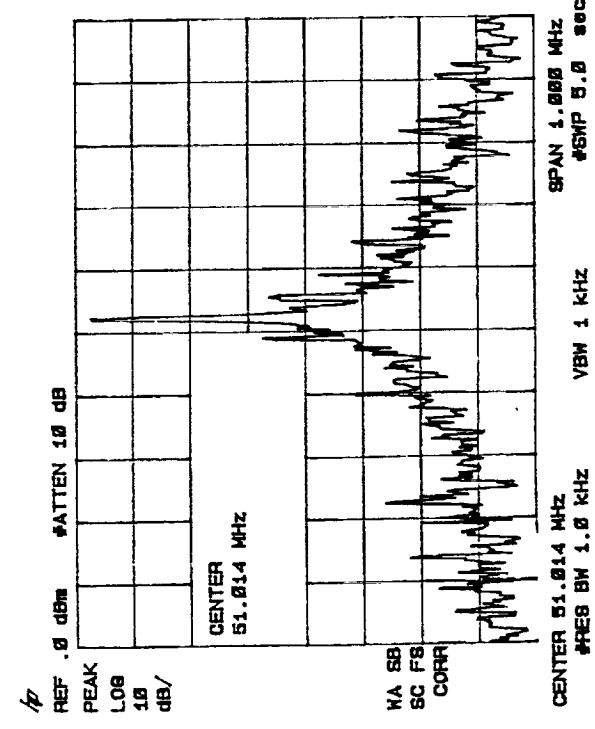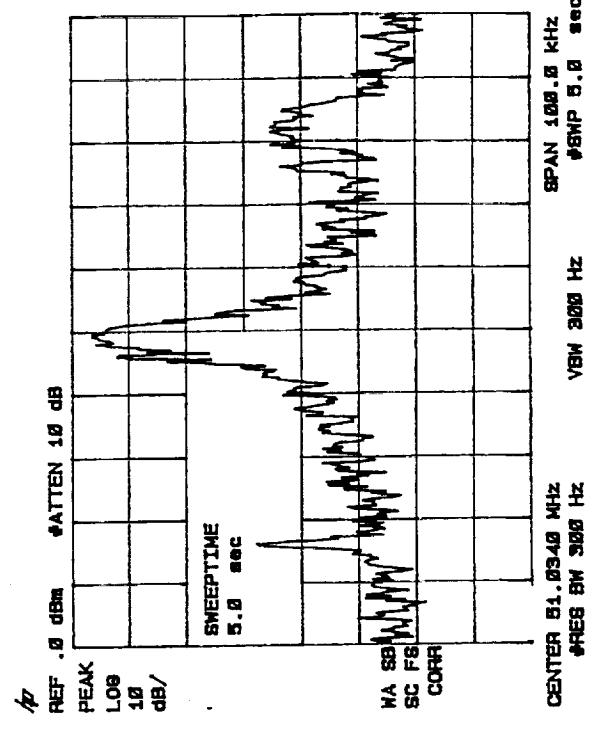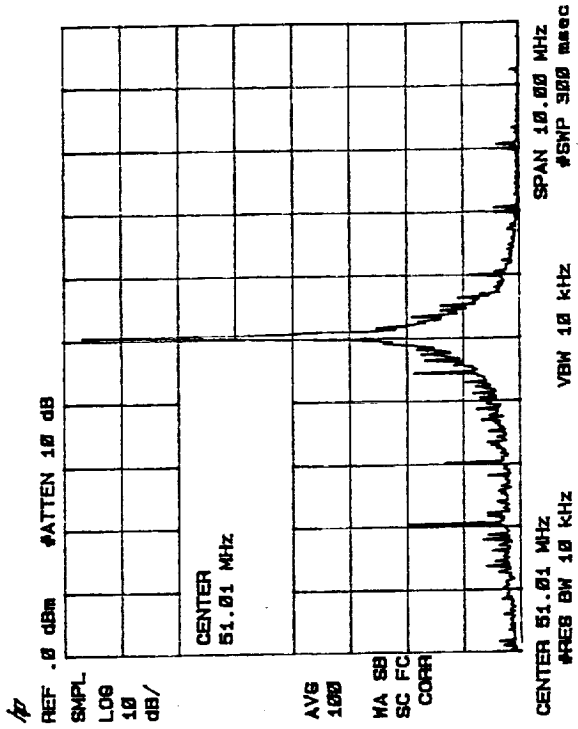
Figure 6. Broadband noise performance of MC1648 LC oscillator, operating open-loop at 50 MHz. Same parameters as Figures 1 and 2.
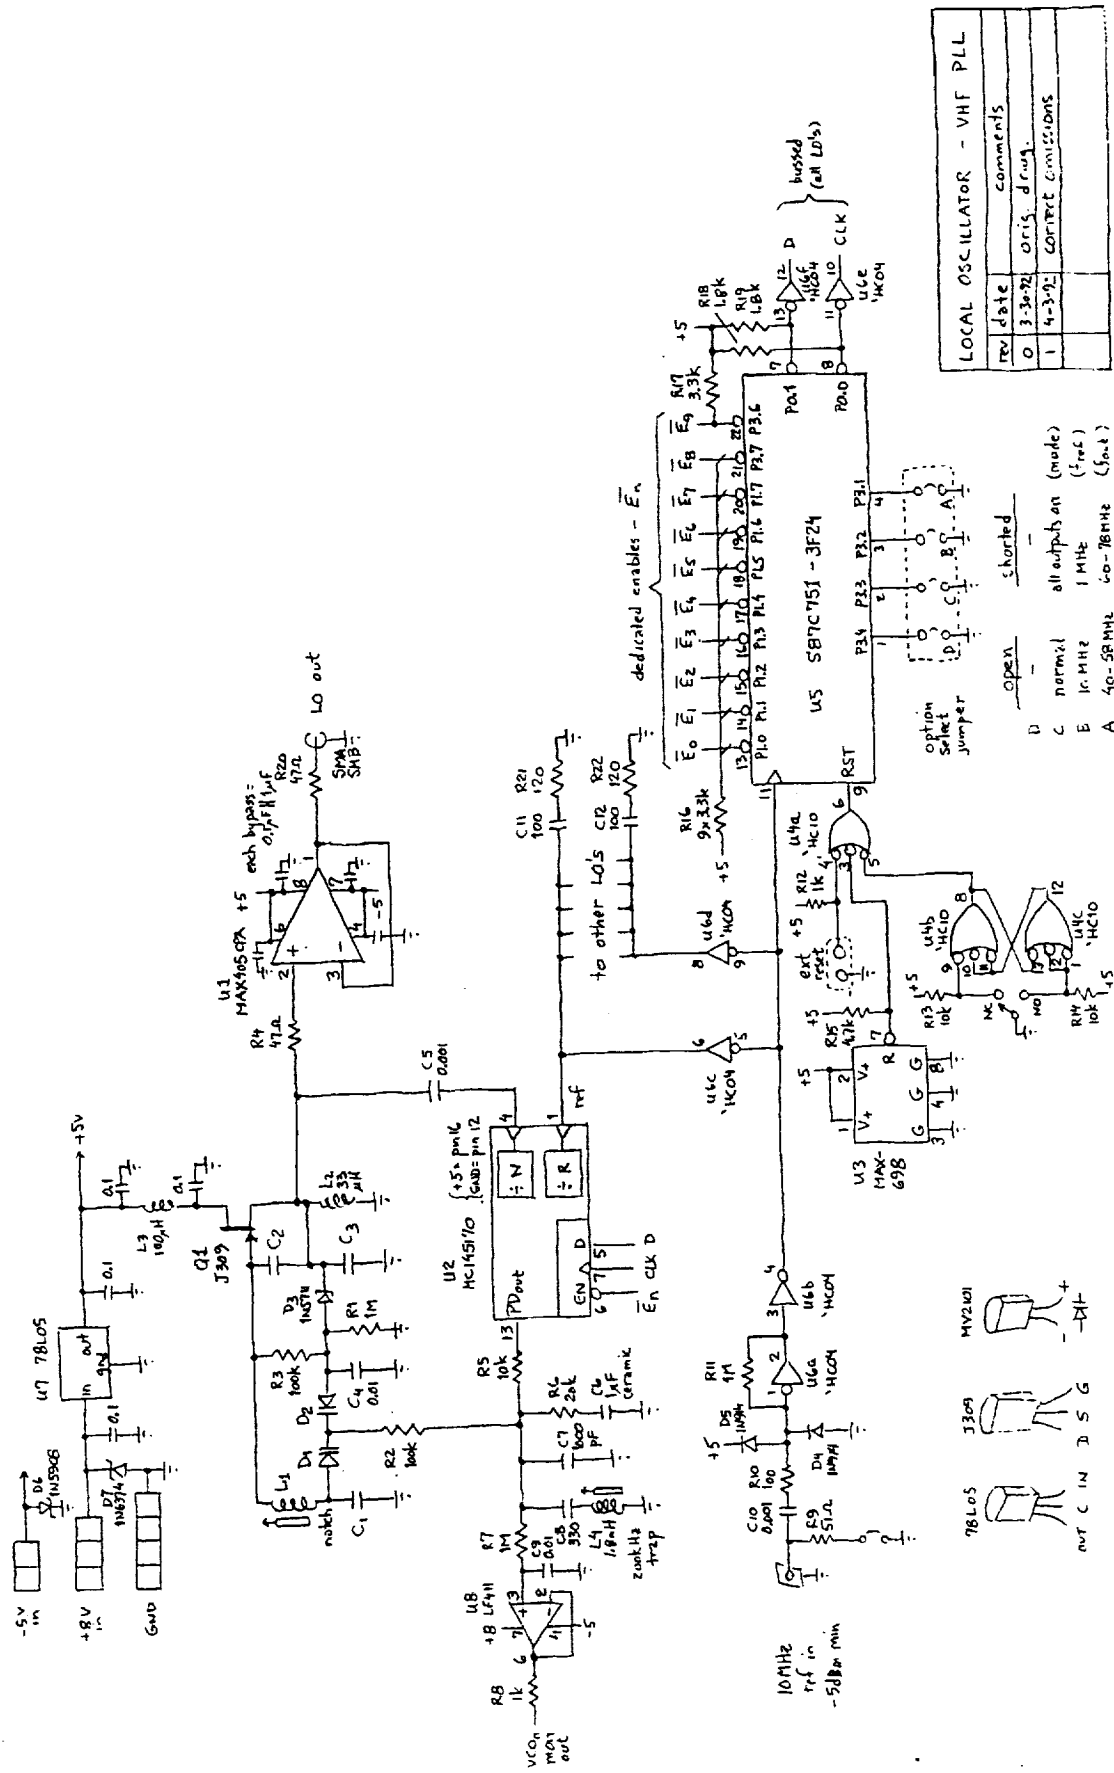
Figure 7. Schematic of 10-channel VHF local-oscillator array, with LC Clapp oscillators locked to 10 MHz reference input. Frequencies programmed via 87C751 microcontroller. Narrow lock range (2 MHz) contributes to spectral purity.
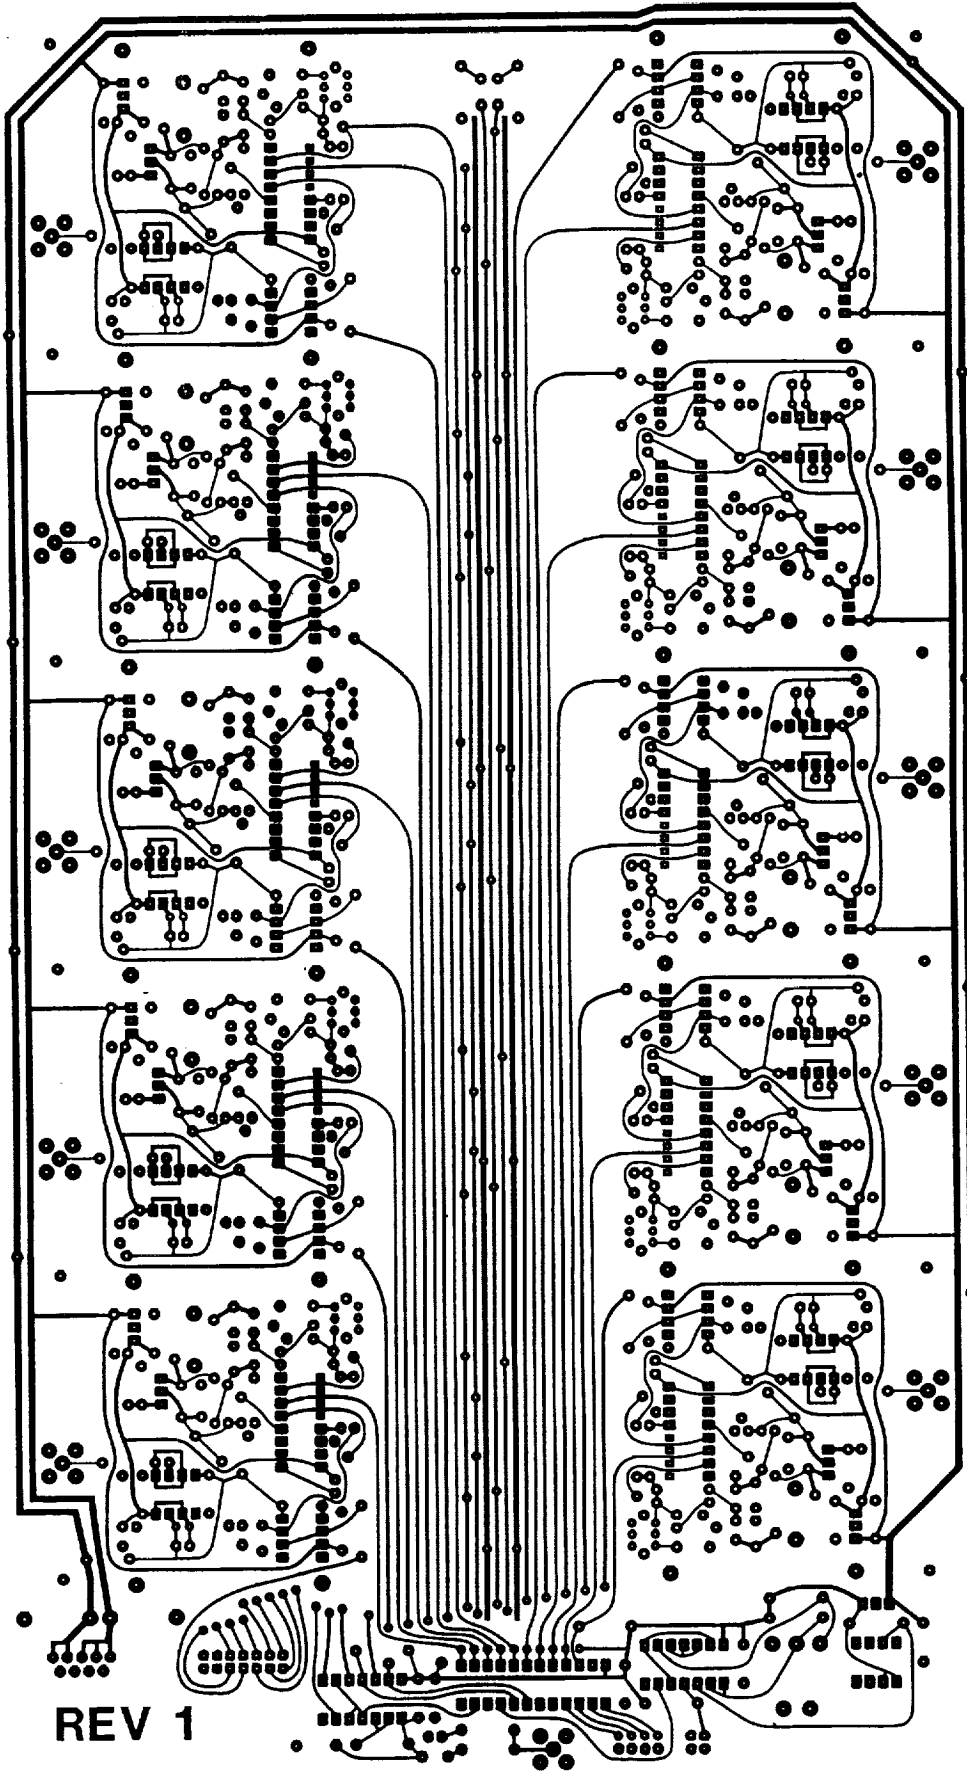
REV 1

Figure 8. Printed circuit pattern (solder side) for local oscillator array. Board size is 7.5 x 14 inches, double sided with ground plane.
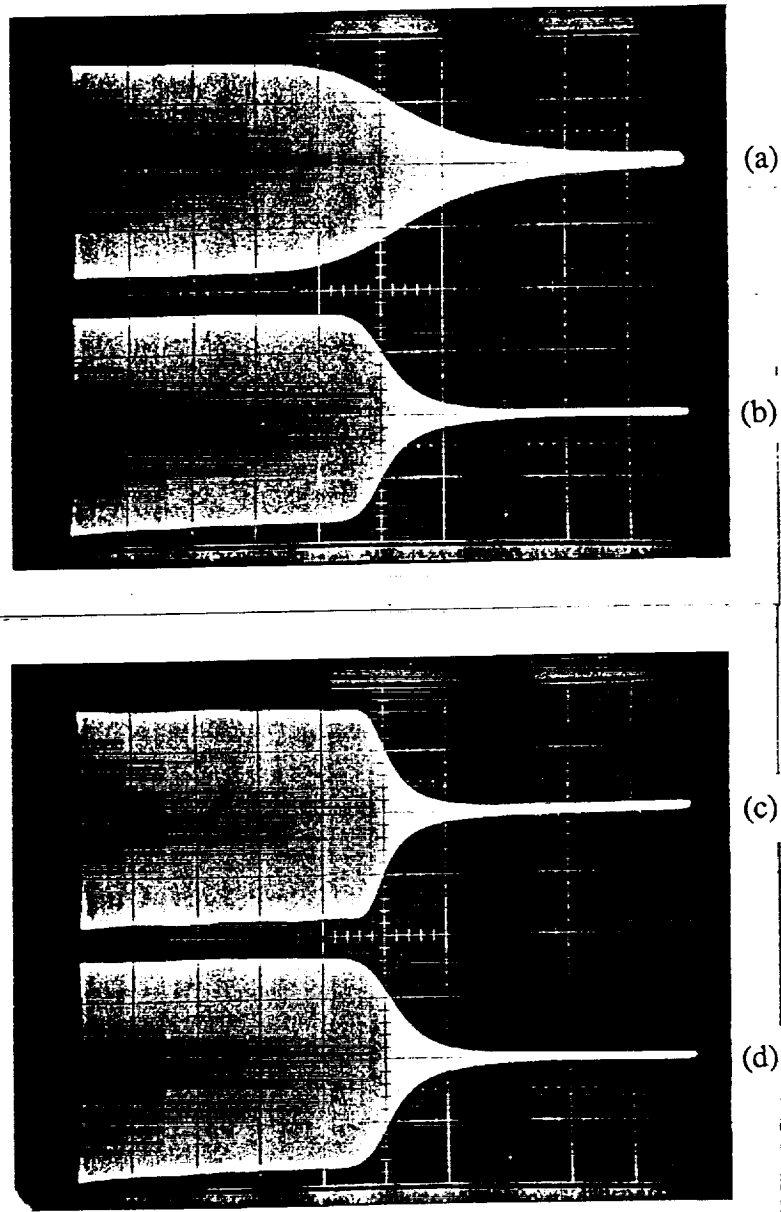
Figure 9. Frequency response of LC anti-alias lowpass filters, all with f_o=1 MHz. Horizontal: linear frequency sweep, 0 to 2 MHz; vertical: linear amplitude scale, constant voltage drive. (a) 5-pole Butterworth; (b) and (d) 7-pole Chebyshev, 0.01dB ripple; (c) 7-pole Chebyshev, 0.1dB ripple.
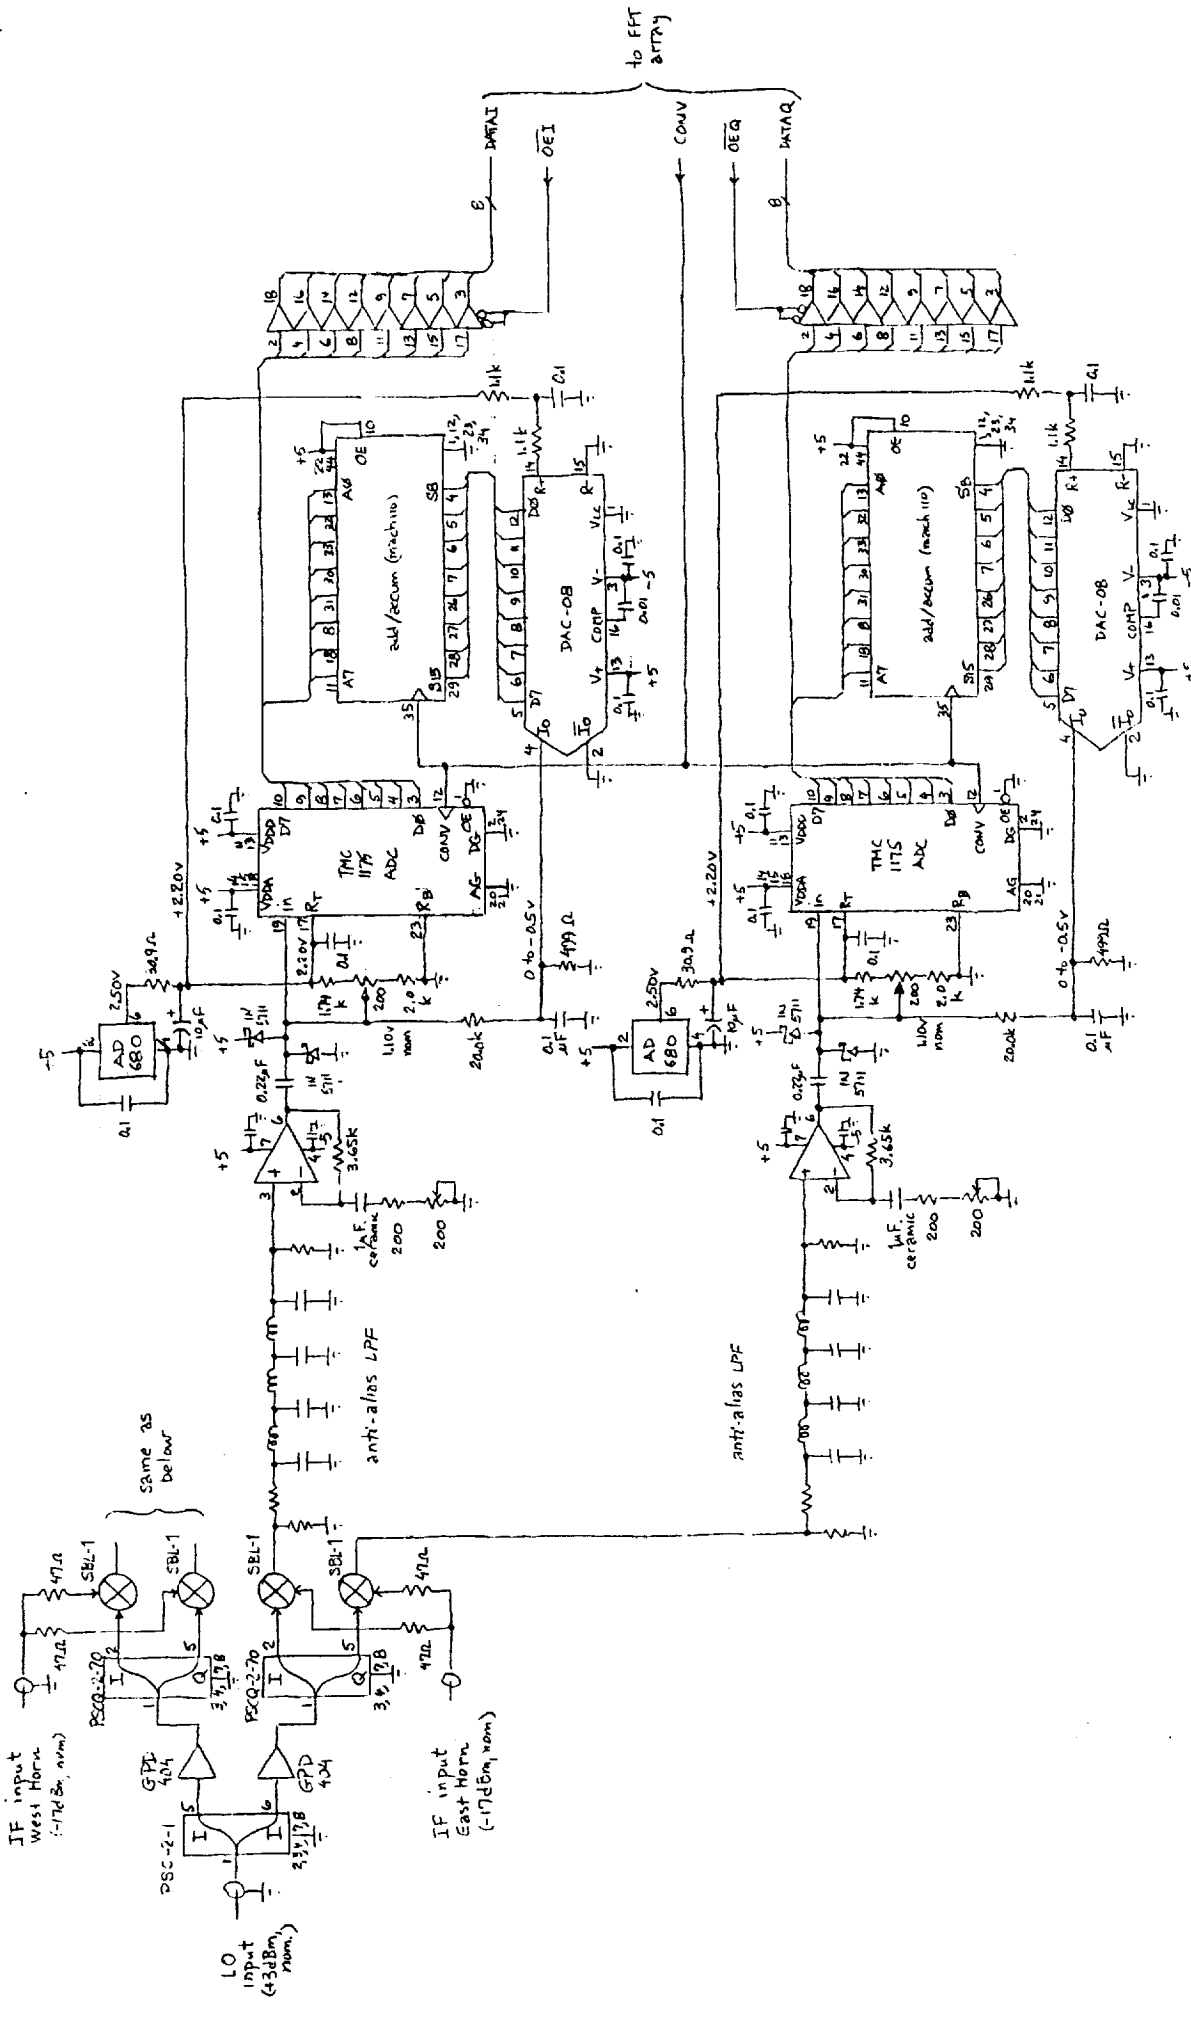
Figure 10. Schematic of quadrature downconverter, with autozeroing feedback to set dc level at ADC input. The Mach 110 is a PAL with 32 macrocells, here implementing a 16-bit full adder with accumulator, clocked at the ADC sampling rate of 2.5 MHz. Loop bandwidth of the autozero circuit is approximately 600 Hz.
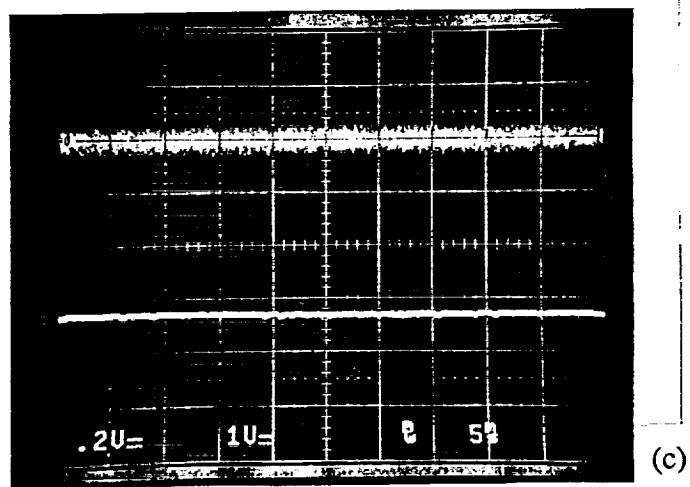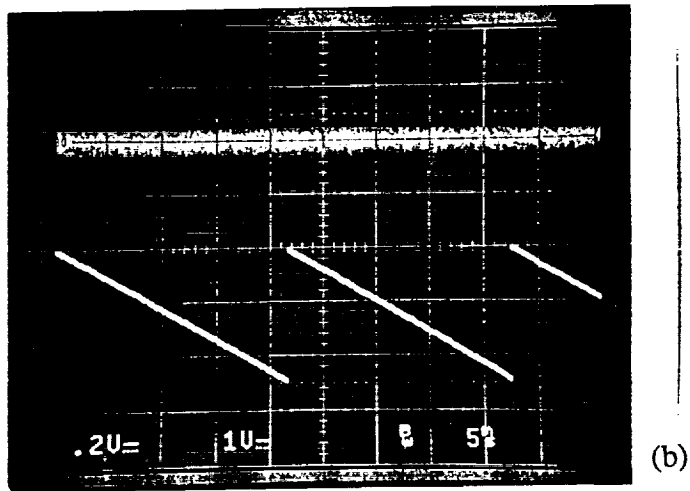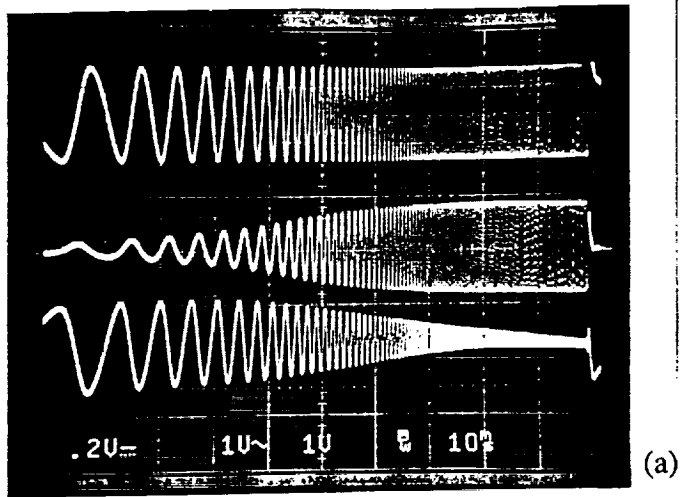
Figure 11. Autozero dynamics, when driven by sinewaves or noise. (a) dc to 10 kHz linear sweep; top trace is constant level input signal, middle trace is signal reaching the ADC, and bottom trace is the autozero feedback. (b) 1 volt broadband (1 MHz) noise, approx 1 volt pp amplitude; top trace in input signal, and bottom trace is the output of the autozero DAC *with feedback loop open*; dc error at input is integrated by the adder-accumulator, generating sawtooth output. (c) same, but with loop closed; note headroom of autozero, whose output can swing 2.5 vertical divisions.
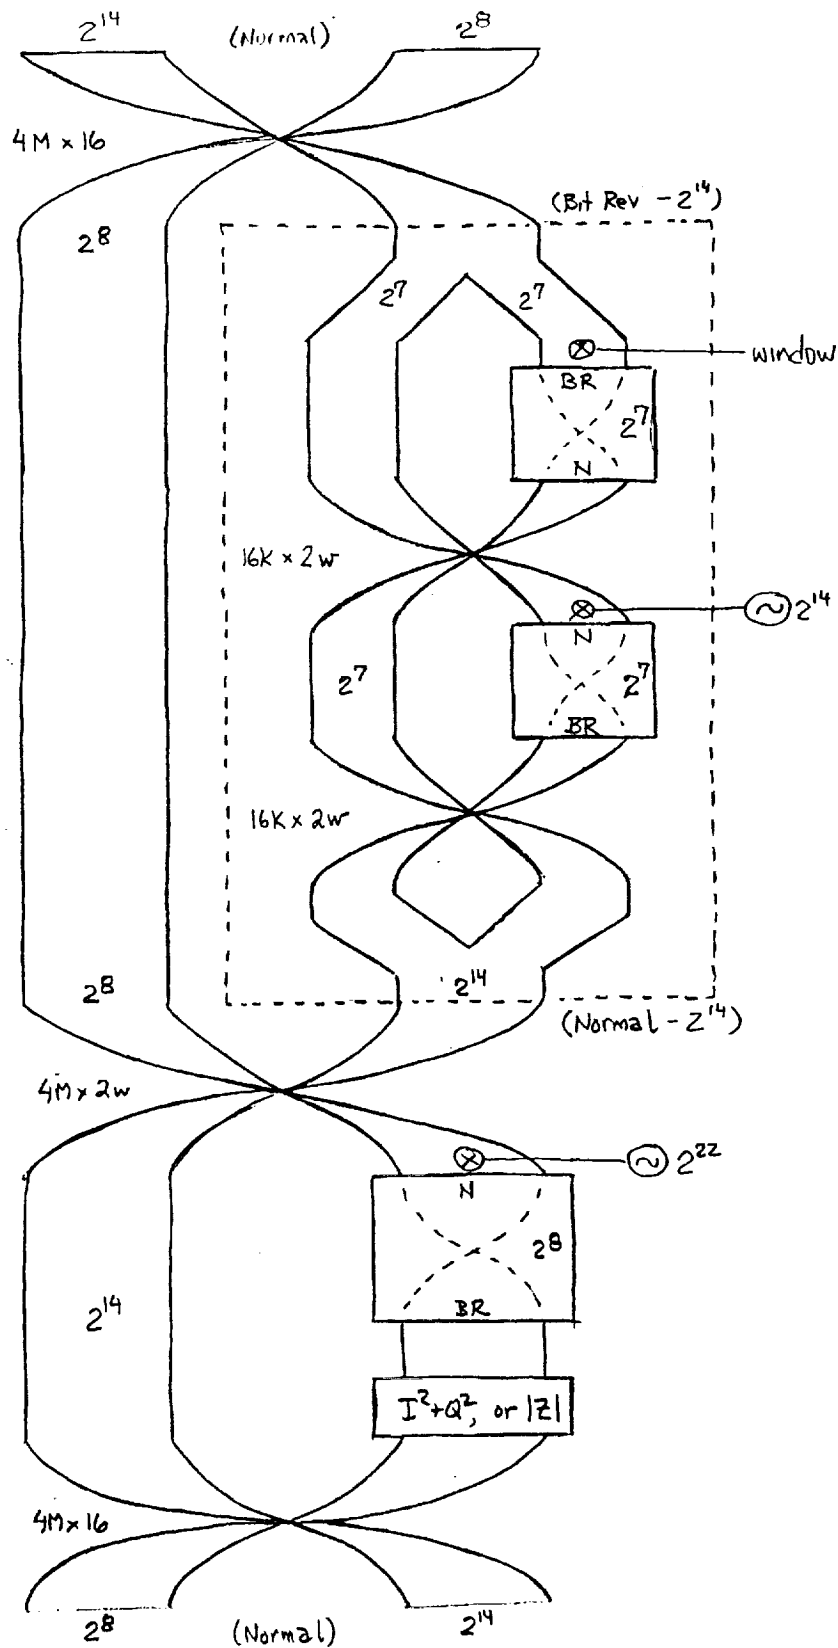
Figure 12. Data flow in a 2^{22}-point FFT, implemented as 256 2^{14}-point FFT's followed by 16384 2^8-point FFT's, with phase rotation by a "twiddle-factor" in between (indicated by a multiplier symbol). The 2^{14}-point FFT is itself a succession of 2^7-point FFT's, again with twiddle factors. The rectangular boxes are FFT's, performed with the Austek A41102 chip; BR means bit-reversed data ordering; N means normally (monotonically) ordered; crossovers represent bit-reverse ordering of the data stream, via "corner-turn" memory buffers. Both twiddle-factor multiplication and input time windowing are performed in the IC's internal complex multiplier.